arXiv:1604.02646v3 [cs.LG] 3 Jan 2017

Visualization Regularizers for Neural Network based Image Recognition

Biswajit Paria
CSE, IIT Kharagpur, India
biswajitsc@iitkgp.ac.in

Vikas Reddy
Mining Engg., IIT Kharagpur, India
vikas.challaram@iitkgp.ac.in

Anirban Santara
CSE, IIT Kharagpur, India
anirban_santara@iitkgp.ac.in

Pabitra Mitra
CSE, IIT Kharagpur, India
pabitra@cse.iitkgp.ernet.in

Abstract

The success of deep neural networks is mostly due their ability to learn meaningful features from the data. Features learned in the hidden layers of deep neural networks trained in computer vision tasks have been shown to be similar to mid-level vision features. We leverage this fact in this work and propose the visualization regularizer for image tasks. The proposed regularization technique enforces smoothness of the features learned by hidden nodes and turns out to be a special case of Tikhonov regularization. We achieve higher classification accuracy as compared to existing regularizers such as the L2 norm regularizer and dropout, on benchmark datasets without changing the training computational complexity.

1 Introduction

Regularization is an important aspect of deep neural network training, to prevent over-fitting in the absence of sufficient data. Usually regularizers restrict the norms of the weight parameters. A commonly used class of regularizers, the Lp norm regularizers, penalize the Lp norms of the weight parameters. Of these, the L1 and L2 norm regularizers are most popular. Other regularizers include soft-weight sharing [1], layer-wise unsupervised pre-training [2], and dropout [3]. It has been shown by Erhan et al. [4] that layer-wise unsupervised pre-training has a regularizing effect during training. There has also been recent works on adaptive dropout [5], which is an improvement over the original dropout. In this paper, we propose a novel regularizer which we call as the visualization regularizer (VR), based on the visual quality of features learned by the hidden nodes. We introduce two variants of the VR regularizer, based on the $\|\cdot\|_1$ and $\|\cdot\|_2$ norms respectively.

Vision tasks benefit from features consisting of primitives recognized by mid-level vision systems. Deep neural networks have been known to learn hierarchical layers of feature representation [6, 7]. On observing the features learned by deep neural networks trained using back propagation, it is seen that in contrast to well defined mid-level features, the node features are often noisy. More meaningful features (smoother features for example) can be favorable to training the network. Our proposed regularizer imposes a constraint on the hidden nodes of a neural network to learn smoother features. Since, the definition of the regularizer depends on the visual property of smoothness, it is only pertinent to domains with a notion of spatial locality, such as images.

We show that the VR regularizer is a special case of Tikhonov regularization [8]. The Tikhonov matrix of the conventional L2 norm regularizer corresponds to an identity matrix multiplied by the regularization weight. Whereas, the Tikhonov matrix for the VR regularizer is more generalised but sparse.

We perform experiments with our regularizer on two benchmark datasets: MNIST [9] and CIFAR-10 [10]. We observe that our regularizer aids in learning and improves the classification accuracy when used alongside other regularizers. However the computational complexity of the VR regularized training algorithm remains the same as that of the unregularized training algorithm.

This paper is organized in five sections. Section II gives a brief introduction to the architecture of deep neural networks and the notation used. The notion of visualization of a node is formally defined in section III. The proposed VR regularizer and the training algorithm along with the complexity analysis are described in section IV. Section V establishes the relationship to Tikhonov regularization. Section VI presents all the experimental results and observations. Finally Section VII concludes with a summary of the achievements and scopes for future work.

2 Deep neural network architecture

The investigations presented in this paper are based on a multi-class classification setting. The notation followed in this paper is as follows: $\mathbf{x}$ denotes the input, W_i and $\mathbf{b}_i$ respectively denote the weights and biases corresponding of the layers, $\mathbf{u}_i$ denotes the pre-activation of the layers, $\mathbf{h}_i$ denotes the activation of layers and $\mathbf{y}$ denotes the output of the neural network. The subscripts i in the notation denotes the ith hidden layer. The following equations describe a deep neural network with l hidden layers.

$$\begin{aligned} \mathbf{u}_1 &= W_1 \cdot \mathbf{x} + \mathbf{b}_1, && (1)\\ \mathbf{h}_i &= g(\mathbf{u}_i),\ \forall 1 \leq i \leq l, && (2)\\ \mathbf{u}_{i+1} &= W_{i+1} \cdot \mathbf{h}_i + \mathbf{b}_{i+1},\ \forall 1 \leq i \leq l-1, && (3)\\ \mathbf{y} &= \mathrm{softmax}(W_{l+1} \cdot \mathbf{h}_l + \mathbf{b}_{l+1}), && (4) \end{aligned}$$

where g, is the activation function, a monotonically increasing non-linear function such as sigmoid, tanh or the rectified linear unit [11].

The loss function used for training is the sum of the classification loss and the regularization term. For a neural network M and a dataset D, the loss function can be written as,

$$L(M, D) = L_c(M, D) + \lambda Reg(M), \tag{5}$$

where $L_c(M, D)$ denotes the classification loss between the output of the neural network and the true class labels, $Reg(M)$ denotes the regularization term and λ denotes the regularizer weight. However, dropout cannot be included in the loss function. It is incorporated into the training algorithm.

3 The notion of visualization of a node

Visualization of a node refers to the visualization of the features learned by the node. Visualization of the nodes of a neural network has been studied by Erhan et al. [12]. They proposed the activation maximization algorithm for visualizing features learned by a node. Following the notion in [12] we define the visualization of a node as follows.

Visualization of a node is defined to be the input pattern(s) which activate(s) the node maximally under the restriction of the L2 norm of the input to be equal to unity. The L2 norm of the input is restricted to unity to prevent the input from becoming unbounded. Formally, the visualization $\mathrm{vis}(n)$ of a node n is defined as,

$$\mathrm{vis}(n) = \underset{\|\mathbf{x}\|_2=1}{\mathrm{argmax}}\, n(\mathbf{x}), \tag{6}$$

where $n(\mathbf{x})$ denotes the activation of node n for input $\mathbf{x}$. Depending on the non-linearity g used in equation 2, solution to the above equation can be unique or multiple. For example, there exists an unique solution for invertible functions like $g = \sigma$ (sigmoid) or $g = \tanh$. Whereas there can be multiple solutions for non-invertible functions such as $g(\cdot) = \max(0, \cdot)$ also known as the rectified linear unit (ReLU). The visualization of an internal node of the neural network can be computed using gradient-ascent as described by Erhan et al. [12] as activation maximization.

Observe that the pre-activation of a node n in the *first* hidden layer for an input vector $\mathbf{x}$ is $\mathbf{w}^\intercal\mathbf{x}$, where $\mathbf{w}$ denotes the weights of the connections coming into the node n. The pre-activation is

Table 1: Visualization loss of example visualizations.

Image (I)	Convolution ($I \otimes K$)	VL(I)
		135.6717
		16.1213
		9.3421
		0.0685

maximized when $\mathbf{x}$ is aligned in the same direction as $\mathbf{w}$ in the appropriate vector space. Since the activation function is a monotonically increasing function, maximization of the pre-activation maximizes the activation too. Consequently, we get the following closed form as one of the possible solutions to equation 6 for nodes n in the first hidden layer.

$$\text{vis}(n) = \frac{\mathbf{w}}{\|\mathbf{w}\|_2} \tag{7}$$

Note that finding a closed form algebraic expression for the visualization of nodes in the higher hidden layers is difficult due to the non-linearity of the activation function.

4 Proposed visualization based regularizer

The VR regularizer is based on the expression of visualization as given in equation 7. We utilize equation 7 and produce an appropriate regularization loss which we include in final training loss. The following subsections give a detailed description of the VR regularizer.

4.1 Smoothness of a visualization

Intuitively, one can determine whether an image is smooth or noisy by looking at the gradients in the image. An image is smooth if it has small gradients. The gradient of an image can be computed by convolving it with a 2D high pass filter. Examples of high pass filters include, first order gradient filters such as the Sobel operator, or second order gradient filters such as the Laplacian operator. Larger the pixel values in the convolution, the larger the gradients in the original image, and greater presence of noise in the image. We utilize this intuition to give a formal definition of smoothness.

Consider a convolution $I \otimes K$ of image I with kernel K, where K is a high pass filter like the laplacian kernel. We define the smoothness of an image to be the negative sum of squares of pixel values of $I \otimes K$. Equivalently, we can define the visualization loss VL of image I as,

$$\begin{aligned} I' &= I \otimes K, \\ \text{VL}(I) &= \sum_{i,j} (I' \cdot I')_{ij}, \end{aligned} \tag{8}$$

where "$\cdot$" denotes the element-wise product, also known as the Schur or Hadamard product. The visualization loss is the negative of the smoothness of an image. Lower the visualization loss, smoother is the image. Table 1 shows the visualization loss for some example visualizations.

4.2 Visualization loss as a regularizer

Classification tasks using deep neural networks benefit from the high-level of abstractions achieved in the higher layers of the neural network. Deep neural networks are intended to utilize low-level pixels to learn mid-level features and finally high-level features. We propose the visualization regularizer (VR) to constrain the nodes in the first hidden layer to learn features with qualities similar to mid-level visual features. This constraint is intended to facilitate the discovery of high-level abstractions more effectively.

Informally, we define the VR regularizer as a regularizer to reduce the visualization loss of the nodes of the neural network. In other words the VR regularizer makes the nodes learn smooth or less noisy features.

The following sub-sections give a more detailed description of the VR regularizer.

4.2.1 Regularizer expression

Let $U(M)$ denote the set of nodes of the first hidden layer of a neural network M and for a node $n \in U(M)$, let $\mathbf{w}_n$ denote the weights of the connections incoming into the node.

From equation 7 we know that visualization of a node in the first hidden layer is proportional to the weights of the connections coming into the node. Hence the visualization loss of the node is proportional to the visualization loss of the weight vector coming into the node. Therefore we can use $\mathrm{VL}(\mathbf{w}_n)$ as a surrogate for the visualization loss of the visualization of the node n. The difficulty of computing an algebraic expression for the visualization of nodes in higher hidden layers, limits the usage of the surrogate to nodes in the first hidden layer only.

We define the visualization loss VL of a neural network M as,

$$\mathrm{VL}(M) = \sum_{n \in U(M)} \mathrm{VL}(\mathbf{w}_n). \tag{9}$$

The network training loss function can thus be defined as,

$$L(M, D) = L_c(M, D) + \mu\, \mathrm{VL}(M) + \lambda L_2'(M), \tag{10}$$

where $L_2'(M)$ denotes the L2 norm regularization term for all weights except the weights coming into the first hidden layer. The remaining notation are as described in section 2.

4.2.2 Gradient of the visualization regularizer

For the visualization loss to be used as a regularizer, its gradient must be computed with respect to the model parameters. Automatic gradient computation libraries such as Tensorflow [13] and Theano [14], obviate the need to compute the gradients manually. However they are an additional overhead on computational resources, and hence for manual computation of gradients for increased efficiency, we derive the expression of the derivative of the VR regularizer. This is also useful in the case of embedded computing, where computational resources are limited. In the following paragraphs we derive the gradient for a general kernel K of size $(2k+1) \times (2k+1)$.

For simplicity in computing the expression of the gradient, we index the elements of the kernel relative to the central element as shown in equation 11. The element at the center is indexed $(0, 0)$. All other elements are indexed according to their position relative to the central element.

$$K = \begin{bmatrix} a_{-k,-k} & \dots & a_{-k,0} & \dots & a_{-k,k} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ a_{0,-k} & \dots & a_{0,0} & \dots & a_{0,k} \\ \vdots & \ddots & \vdots & \ddots & \vdots \\ a_{k,-k} & \dots & a_{k,0} & \dots & a_{k,k} \end{bmatrix} \tag{11}$$

Let N_K denote the set of indices of the kernel matrix.

$$N_K = \{(i, j) \mid i, j \in \{-k, \dots, k\}\} \tag{12}$$

Consider an image I with dimensions n, m. Let $S(i,j)$, corresponding to the (i,j)th pixel of I, be defined as follows.

$$\begin{aligned} S(i,j) &= \{(r,(p,q)) \mid (p,q) = (i,j) + r,\ r \in N_K \\ &\quad \text{and } 0 \leq p < n,\ 0 \leq q < m\} \end{aligned} \tag{13}$$

Informally, $S(i,j)$ contains the set of valid indices (p,q), along with their position r relative to (i,j), that need to be considered while computing the convolution for the (i,j)th pixel.

A full convolution of the image I can be described as

$$I' = I \otimes K = \left[\sum_{(r,(p,q)) \in S(i,j)} a_r I_{pq} \right]_{ij}. \tag{14}$$

It follows from the definition that a pixel I_{pq} present at a position r, relative to (i,j), has the coefficient a_r in I'_{ij}.

The visualization loss is

$$\mathrm{VL}(I) = \sum_{i,j} {I'}^2_{ij}. \tag{15}$$

The partial derivative of the visualization loss with respect to a pixel I_{ij} of the image I is

$$\frac{\partial}{\partial I_{ij}} \mathrm{VL}(I) = \sum_{p,q} \frac{\partial}{\partial I_{ij}} {I'}^2_{pq}. \tag{16}$$

Observe that in the above equation I_{ij} occurs in I'_{pq} only for (p,q) such that $(r,(p,q)) \in S(i,j)$. Moreover, if (p,q) is present at position r relative to (i,j), then (i,j) is present at position $-r$ relative to (p,q). It follows that I_{ij} has the coefficient a_{-r} in I'_{pq}. Hence the derivative can be computed as,

$$\begin{aligned} \frac{\partial}{\partial I_{ij}} \mathrm{VL}(I) &= \sum_{(r,(p,q)) \in S(i,j)} \frac{\partial}{\partial I_{ij}} {I'}^2_{pq} \\ &= \sum_{(r,(p,q)) \in S(i,j)} 2 I'_{pq} \frac{\partial}{\partial I_{ij}} I'_{pq} \\ &= \sum_{(r,(p,q)) \in S(i,j)} 2 I'_{pq} a_{-r}. \end{aligned}$$

Further, using equation 14, we can write,

$$\frac{\partial}{\partial I} \mathrm{VL}(I) = 2(I' \otimes K^I) = 2((I \otimes K) \otimes K^I), \tag{17}$$

where K^I denotes the kernel matrix formed by flipping K both horizontally and vertically. Formally,

$$K^I_{ij} = K_{(-i,-j)}. \tag{18}$$

Equation 17 allows us to compute the gradient efficiently and in a scalable manner. All popular GPU programming frameworks provide libraries for scalable convolutions. Figure 1 illustrates the gradient computation for the visualization loss.

4.2.3 Regularized training algorithm

Training requires computing the gradient for the regularized loss function with respect to the parameters of the network. As evident from equation 10, the gradient of the loss function can be computed by first computing the gradients of $L_c(M,D)$, $L'_2(M)$ and $\mathrm{VL}(M)$, and then computing their sum. The gradients of $L_c(M,D)$ and $L'_2(M)$ can be computed using back-propagation and partial derivatives respectively.

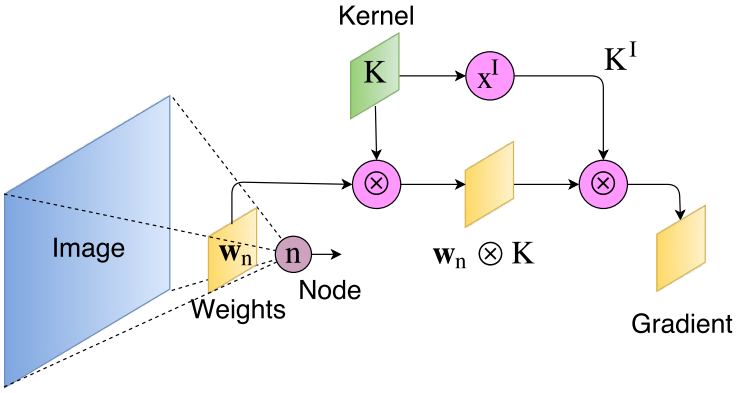

Figure 1: Computation of gradient for visualization loss. $\otimes$ denotes convolution.

By equation 9, the gradient of $\mathrm{VL}(M)$ is the sum of gradients of $\mathrm{VL}(\mathbf{w}_n)$ for $n \in U(M)$. Note that the gradient $\frac{\partial}{\partial w}\mathrm{VL}(\mathbf{w}_n)$ is zero if $w \notin \mathbf{w}_n$. In other words, the gradient is zero if w does not belong to the set of weights incoming to node n. Hence we only need to compute the gradients $\frac{\partial}{\partial \mathbf{w}_n}\mathrm{VL}(\mathbf{w}_n)$ for all $n \in U(M)$. These can be computed using equation 17. The full algorithm described in figure 2. The algorithm can be extended to using dropout and momentum.

Computing each $\mathbf{z}_n$ in figure 2 takes $O(|\mathbf{w}_n|)$. Consequently, computing the gradient of the VR regularizer takes $O(|\mathbf{w}|)$ time, which is of the same order as computing the gradient of the L2 norm regularizer. Thus, the VR regularizer does not impose additional overhead in the computational complexity per iteration.

4.3 A variant of the VR regularizer

In our description of the VR regularizer we have defined the smoothness of an image I as the sum of squares of the convoluted image I' (equation 8), which is actually the square of the 2-norm ($\|\cdot\|_2$) of the flattened vector of I'. Analogously, we can also define an 1-norm variant of the VR regularizer. To distinguish between the 1-norm and 2-norm variants, we denote the respective losses by VL_1 and VL_2. Similar to the definition of VL_2 in equation 8, VL_1 can be defined as follows.

$$\begin{aligned} I' &= I \otimes K, \\ \mathrm{VL}_1(I) &= \sum_{i,j} |I'_{ij}|. \end{aligned} \tag{19}$$

Thus, the training loss can be modified as,

$$L(M, D) = L_c(M, D) + \mu_1 \mathrm{VL}_1(M) + \mu_2 \mathrm{VL}_2(M) + \lambda L'_2(M), \tag{20}$$

Similar to the expression of the gradient for VL_2 (equation 17), we can derive the gradient of VL_1 as

$$\frac{\partial}{\partial I}\mathrm{VL}_1(I) = 2(\mathrm{sign}(I') \otimes K^I) = 2(\mathrm{sign}(I \otimes K) \otimes K^I), \tag{21}$$

where $\mathrm{sign}(M)$ denotes the elementwise application of the signum[1] function on the elements of the matrix M.

[1] $\mathrm{sign}(x) = \begin{cases} 1, & \text{if } x > 0 \\ 0, & \text{if } x = 0 \\ -1, & \text{if } x < 0 \end{cases}$

procedure TRAIN($D, M, K, \mu, \lambda, \alpha$)
    ▷ D: the training data, M: the model respectively.
    ▷ K: Kernel for the VR regularizer.
    ▷ μ: VR regularizer weight, λ: L2 regularizer weight.
    ▷ α: learning rate.
    ▷ Returns trained model M

    ▷ Get the weight variables from the model
    $\mathbf{w} := w(M)$
    Initialize $\mathbf{w}$

    while does not converge **do**
        for minibatch sample D_m **do**
            ▷ Step 1: Compute using back-propagation.
            $\mathbf{u} \leftarrow \nabla_{\mathbf{w}} L_c(D_m, M)$
            ▷ Step 2: Gradient for L2 regularizer.
            $\mathbf{v} \leftarrow \nabla_{\mathbf{w}} L_2'(M)$

            ▷ Step 3: Gradient for VR regularizer.
            ▷ Iterate over the nodes in the 1st fully connected layer.
            for $n \in U(M)$ **do**
                ▷ Compute $\mathbf{z}_n = \frac{\partial}{\partial \mathbf{w}_n} \mathrm{VL}(\mathbf{w}_n)$
                ▷ using equation 17
                $\mathbf{z}_n = 2((\mathbf{w}_n \otimes K) \otimes K^I)$
            end for
            ▷ Concatenate the gradients.
            $\mathbf{z} = (\dots, \mathbf{z}_n, \dots)$

            ▷ Step 4: Compute total gradient.
            $\mathbf{g} \leftarrow \mathbf{u} + \lambda \mathbf{v} + \mu \mathbf{z}$
            ▷ Step 5: Take step.
            $\mathbf{w} \leftarrow \mathbf{w} - \alpha \mathbf{g}$
        end for
    end while
    return M
end procedure

Figure 2: Regularized training algorithm

5 Relationship with Tikhonov regularization

Tikhonov regularization was originally developed for solutions to ill-posed problems [8]. For example, L2 regularization, a special case of Tikhonov regularization is used to compute solutions to regression problems for which rank deficient matrices are encountered while computing their solutions. The solution to regularized least squares regression is given by,

$$w = \underset{w}{\operatorname{argmin}} \|Xw - y\|^2 + \lambda \|\Gamma w\|^2, \tag{22}$$

where Γ is the Tikhonov matrix. The L2 regularizer corresponds to $\Gamma = I$.

We show that the VL_2 regularizer is also a special case of Tikhonov regularization. Let $\mathbf{w}$ be the concatenation of all the weights in $\mathbf{w}_n$ for $n \in U(M)$. From equation 15, we can see that $\mathrm{VL}(\mathbf{w}_n)$ is the sum of squared terms of the form ${\mathbf{w}'_{ij}}^2 = (\sum_t c_t w_t)^2$, where $w_t \in \mathbf{w}$, $c_t \in K$ and $\mathbf{w}' = \mathbf{w}_n \otimes K$. For example $\mathbf{w}'_{ij} = (w_p - w_q - w_r)^2$ can be represented in the given form where $c_p = 1, c_q = -1, c_r = -1$, and all other c_t are zero.

The Tikhonov matrix Γ can be constructed as follows. Consider the expression of $\mathrm{VL}(M) = \sum_{n \in U(M)} \mathrm{VL}(\mathbf{w}_n)$ consisting of the sum of p such squared terms. Let the ith term in this expression be $s_i = (\sum_t c_{it} w_t)^2$. Then, Γ and consequently $\mathrm{VL}(M)$ in terms of Γ are, respectively,

$$\Gamma = [c_{it}]\,, \text{ and } \mathrm{VL}(M) = \|\Gamma \mathbf{w}\|^2. \tag{23}$$

In practice, it can be assumed that the kernel matrix K is has a constant size. Hence it follows, that the expression of s_i consists of only a constant number of non-zero c_{it} since the number of non-zero c_{it} is bounded by the size of the kernel matrix K. Hence, the number of non-zero entries in Γ is $O(p \cdot \mathrm{size}(K)) = O(p)$, whereas the total number of entries in Γ is $p|\mathbf{w}|$, concluding the sparsity of Γ.

6 Experiments and observations

We experimented on the MNIST [9] and CIFAR-10 [10] datasets and compared the classification accuracy of our algorithm using the VR regularizers with other regularizers.[2]

6.1 Experimental setting

For classification of MNIST digits, we experimented with fully connected architectures and convolutional architectures, and we only experimented with convolutional architectures for the classification of CIFAR-10 objects. For fully connected architectures, we applied the VR regularizer to the weights of layer immediately after the input layer i.e. the first hidden layer. For convolutions however, we applied the VR regularizer to the weights of the fully connected layer immediately after the last convolutional layer. The detailed layerwise descriptions of the architectures used in our experiments are given in table 2.

We used the laplacian kernel for the VR regularizer defined as follows.

$$K = \begin{bmatrix} -1 & -1 & -1 \\ -1 & 8 & -1 \\ -1 & -1 & -1 \end{bmatrix} \tag{24}$$

We used the mean cross-entropy loss over the mini-batches as the classification loss. The total training loss is as defined in equation 20. We trained the neural network using stochastic gradient descent with momentum [15].

The the initial learning rate for all the models was fixed to 0.01. The model was trained for 2300 epochs and 250 epochs for CIFAR-10 and MNIST respectively. For CIFAR-10, the learning rate

[2]The experiment code in Theano is available at `https://github.com/biswajitsc/VisRegDL`. An implementation in Tensorflow is also available at `https://github.com/cvikasreddy/VisReg`.

Table 2: Description of the network architectures used in our experiments. fc(n) denotes a fully connected layer with n nodes, dropout(p) denotes a dropout layer with nodes dropped with probability p, conv($s \times s$, c) denotes a convolutional layer with a kernel of size $s \times s$ and c output channels, and maxpool($s \times s$) denotes a maxpooling layer with a $s \times s$ sized kernel.

Dataset / Architecture	Description
MNIST / fully connected	input(784) – fc(1000) – dropout(0.3) – fc(1000) – dropout(0.3) – fc(1000) – output(10)
MNIST / convolutional	input(28×28) – conv(3×3, 64) – conv(3×3, 64) – dropout(0.1) – maxpool(3×3) – dropout(0.1) – fc(1024) – output(10)
CIFAR-10 / convolutional	input(32×32) – conv(5×5, 64) – dropout(0.1) – maxpool(3×3) – conv(5×5, 64) – dropout(0.1) – maxpool(3×3) – conv(5×5, 64) – dropout(0.1) – maxpool(3,3) – dropout(0.1) – fc(384) – dropout(0.1) – fc(192) – dropout(0.1) – out(10)

Table 3: Accuracies for various regularizer settings. '-' denotes that the corresponding regularizer was not used. VL_1 and VL_2 denote the respective mean losses over the whole dataset.

Dataset / Architecture	μ_1	μ_2	λ	Acc. %
MNIST / fully connected	-	-	-	98.61
	-	-	0.02	98.75
	-	0.01	0.01	**98.81**
	0.01	-	-	98.39
MNIST / convolutional	-	-	-	99.21
	-	-	0.01	99.16
	-	0.02	-	99.27
	0.01	-	-	**99.29**
CIFAR-10 / convolutional	-	-	-	80.49
	-	0.01	0.004	81.91
	0.01	-	0.004	82.12
	-	0.001	0.004	82.09
	0.001	-	0.004	**82.18**

was reduced by a factor of 1.3 every 500 epochs. For MNIST, it was reduced to 0.005 after the 75th epoch and reduced by a factor of 1.3 thereafter every 25th epoch. For finding the optimal values of the regularization weights, we performed a randomized hyper-parameter search with manual fine-tuning.

6.2 Discussion

We compared training various neural network models with combinations of L2, VR, and dropout regularizers. The accuracy and optimal hyper-parameters for various regularizer settings are given in table 3. The parameters μ_1, μ_2, λ and α respectively denote the VR regularizer weights, the L2 regularizer weight, and the learning rate.

From the table it is observed that the VR regularizer leads to an improvement in the classification accuracy. This is observed both for fully connected and convolutional neural networks. Moreover

it is also observed that VR_1 is a better regularizer for convolutional neural networks compared to VR_2, whereas VR_2 works better in the case of fully connected neural networks.

7 Conclusion and future work

In this paper we introduced a new regularizer for deep neural networks trained for image tasks. We formulated the regularizer based on the notion of smoothness of a visualization, and also derived the expression of its gradient. We experimentally observe that the VR regularizer aids in learning and leads to an improvement in the classification accuracy. The VR regularizer introduces a new class of regularizers based on domain assumptions (assumption of smoothness of natural images in our case). Such kind of regularizers can also be used in other domains such as audio and video, where similar to natural images, transitions along the dimensions of the data are not noisy, but mostly smooth.

We can conclude by saying that VR regularizers are a promising direction towards more general regularization techniques which use domain knowledge.

References

[1] Steven J Nowlan and Geoffrey E Hinton. Simplifying neural networks by soft weight-sharing. *Neural computation*, 4(4):473–493, 1992.

[2] Yoshua Bengio, Pascal Lamblin, Dan Popovici, and Hugo Larochelle. Greedy layer-wise training of deep networks. In *In NIPS*, 2007.

[3] Nitish Srivastava, Geoffrey Hinton, Alex Krizhevsky, Ilya Sutskever, and Ruslan Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting. *The Journal of Machine Learning Research*, 15(1):1929–1958, 2014.

[4] Dumitru Erhan, Yoshua Bengio, Aaron Courville, Pierre-Antoine Manzagol, Pascal Vincent, and Samy Bengio. Why does unsupervised pre-training help deep learning? *The Journal of Machine Learning Research*, 11:625–660, 2010.

[5] Jimmy Ba and Brendan Frey. Adaptive dropout for training deep neural networks. In *Advances in Neural Information Processing Systems*, pages 3084–3092, 2013.

[6] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. Deep learning. *Nature*, 521(7553):436–444, 2015.

[7] Yoshua Bengio. Learning deep architectures for ai. *Foundations and trends® in Machine Learning*, 2(1):1–127, 2009.

[8] AN Tikhonov and VY Arsenin. *Solutions of ill-posed problems*. VH Winston and Sons, 1977.

[9] Yann LeCun and Corinna Cortes. The MNIST database of handwritten digits. 1998. URL `http://yann.lecun.com/exdb/mnist`.

[10] Alex Krizhevsky. Learning multiple layers of features from tiny images. *Technical Report, Univ. Toronto*, 2009.

[11] Xavier Glorot, Antoine Bordes, and Yoshua Bengio. Deep sparse rectifier neural networks. In *Aistats*, volume 15, page 275, 2011.

[12] Dumitru Erhan, Yoshua Bengio, Aaron Courville, and Pascal Vincent. Visualizing higher-layer features of a deep network. *Technical Report, Univ. Montreal*, 2009.

[13] Martín Abadi, Ashish Agarwal, Paul Barham, Eugene Brevdo, Zhifeng Chen, Craig Citro, Greg S. Corrado, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Ian Goodfellow, Andrew Harp, Geoffrey Irving, Michael Isard, Yangqing Jia, Rafal Jozefowicz, Lukasz Kaiser, Manjunath Kudlur, Josh Levenberg, Dan Mané, Rajat Monga, Sherry Moore, Derek Murray, Chris Olah, Mike Schuster, Jonathon Shlens, Benoit Steiner, Ilya Sutskever, Kunal Talwar, Paul Tucker, Vincent Vanhoucke, Vijay Vasudevan, Fernanda Viégas, Oriol Vinyals, Pete Warden, Martin Wattenberg, Martin Wicke, Yuan Yu, and Xiaoqiang Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. URL `http://tensorflow.org/`. Software available from tensorflow.org.

[14] Theano Development Team. Theano: A Python framework for fast computation of mathematical expressions. *arXiv e-prints*, abs/1605.02688, May 2016. URL http://arxiv.org/abs/1605.02688.

[15] Y Nesterov. A method of solving a convex programming problem with convergence rate {O}(1/kˆ2). *Soviet Mathematics Doklady*, 27(2):372–376, 1983. URL http://www.core.ucl.ac.be/{~}nesterov/Research/Papers/DAN83.pdf.